

PATENT
70111.00008

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: GARY SHUSTER

Serial No.: 09/932,431

Filed: August 17, 2001

Title: METHOD AND APPARATUS FOR
IMPROVING BANDWIDTH EFFICIENCY IN A
COMPUTER NETWORK

REPLY TO SUPPLEMENTAL EXAMINER'S ANSWER

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir or Madam:

The appellant filed a Second Amended Appeal Brief (the "Appeal Brief") on September 1, 2006 in the above-identified application, to which the Patent and Trademark Office (the "Office") provided an Examiner's Answer on December 4, 2006. The appellant filed a Reply to Examiner's Answer (the "First Reply") on February 5, 2007, which is the first business day following February 4, 2007. The Office mailed a second or Supplemental Examiner's Answer on February 8, 2007 (the "Second Answer"), which is substantively as the first Examiner's Answer, except for responding to the appendices of the Appeal Brief for which a response was omitted in the first Answer. The appellant respectfully submits this Reply to the Examiner's Supplemental

Answer (the “Second Reply”) on the first business day following April 8, 2007, pursuant to 37 CFR 41.41.

With respect to claims 21-36 currently at issue, the cited reference (Rakoshitz) fails to disclose “each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference”. *Verdegaal Bros. v. Union Oil Co. of Calif.*, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Rakoshitz therefore does not anticipate claims 21-36 under 35 U.S.C. § 102(e), as set forth in the Appeal Brief. The issues as organized in the Second Answer are addressed in turn below.

I. REJECTION UNDER 35 U.S.C. § 102(e) OF CLAIMS 21-23, 26, 28-31, 34 AND 36

Issue 1: Rakoshitz discloses software designed to operate at a gateway point, and fails to disclose operating software at both communication links and file storage locations, as defined by claims 21 and 29.

The Second Answer cites Rakoshitz at col. 9:27-38, which states that a bandwidth management tool “can be deployed at any appropriate point in the network data path.” Specifically, Rakoshitz also notes, at col. 8:45-55, that the bandwidth management tool can be deployed at:

1. a network’s Internet access link,
2. a private WAN link to a remote corporate site,
3. an access to a server farm, or

4. key servers.

The Second Answer argues that these citations prove that the “bandwidth management tool can be deployed at multiple locations in a variety of network scenarios.” Second Answer, p. 7:13-15. This somewhat vague statement misses the point, at best. At worst, it stretches the disclosure of Rakoshitz to include subject matter that is plainly absent.

If the statement means merely that Rakoshitz discloses that the **same bandwidth management tool** might be deployed at different points in a network, it misses the point. In that case, it is unrebutted that, as stated on pages 9-10 of the Appeal Brief, Rakoshitz “does not disclose distributing different interoperable bandwidth management modules at different network locations,” and therefore does not disclose the elements defined by claims 21 and 29.

If the statement in the Answer is means that Rakoshitz discloses distributing different functions of a bandwidth management tool to different locations in a network, it is plainly incorrect. Rakoshitz teaches that the entire tool is deployed at each network point. On page nine of the Appeal Brief, no less than nine excerpts from Rakoshitz are provided, each proving an example of deploying the entire bandwidth management tool at a single point. In fact, **every exemplary embodiment** of Rakoshitz discloses implementing the bandwidth management tool at an access point or gateway. Col.

11:5:61. Rakoshitz even touts deployment at “a single point of access” as one of the advantages of the invention. Col. 3:16-23.

Indeed, it would not make sense to install a bandwidth management tool as disclosed by Rakoshitz at any location other than an access point or gateway. The bandwidth management tool of Rakoshitz uses a “FAIR module” that “generally implement[s] traffic control and manages bandwidth of incoming and outgoing information to and from the network or link.” Col. 12:35-39. Rakoshitz further discloses that “Flow Analysis and Intelligent Regulation (“FAIR”) implements traffic control based on a combination of flow control and queuing algorithms.” Col. 12:39-42. Here, Rakoshitz discloses managing incoming and outgoing information using flow control and queuing. Indeed, Rakoshitz consistently discloses implementing traffic control based on an analysis of data flowing through a network node, and not based on stored files. Col. 13:13-15. The operation of Rakoshitz’s traffic control method is described in detail in connection with Fig. 8, where Rakoshitz discloses that “[i]n general, a flow of information or data or packets of information enter **a gateway point, where the present tool sits**.” Col. 15: 58-60 (emphasis added).

Thus, Rakoshitz fails to disclose, either expressly or inherently, the combination of steps defined by claim 21, namely:

monitoring bandwidth usage of a communication link for connecting a server group to a wide area network, using software operably associated with the communication link;

distributing a rule set to individual servers of the server group, wherein the rule set defines rules for limiting serving of data from the individual servers depending on file type and a current state of the bandwidth usage;

characterizing files stored in operable association with the individual servers according to type, using software operating on the individual servers;

informing the individual servers of the current state of the bandwidth usage as monitored by the software operably associated with the communication link.

Similar limitations are defined in claim 29. Claims 21 and 29 define a method or system in which different functions are performed at different locations; for example, monitoring of bandwidth usage occurs at a communication link, while characterization of stored files occurs at individual servers of a server farm. Rakoshitz does not disclose distributing functionality of a bandwidth management tool in the claimed manner, or in any manner save for implementing the entire tool at a point or points of a network.

Issue 2: Rakoshitz does not disclose the claimed elements of claims 21 and 29, including characterization of stored files separately from bandwidth monitoring.

The Second Answer cites col. 2:66-67 and 3:1-15 to support an argument that Rakoshitz discloses separate software performing separate functions. This argument merely knocks down a straw man. Rakoshitz does not disclose what is claimed, which is characterization of stored files at a location separate from where bandwidth is monitored, or at any other location. At 2:66-67 and 3:1-15, Rakoshitz discloses:

In still an alternative embodiment, the present invention provides a novel bandwidth-profiling tool. The present bandwidth profiling tool

includes a variety of computer codes to form computer software or a computer program, which is stored in computer memory. The program includes a first code that is directed to measuring a data rate for a flow of information from an incoming source, which is coupled to a network of computers. The program also has a second code that is directed to categorizing the data rate from the flow of information based upon at least one of a plurality of traffic classes and a third code that is directed to outputting a visual representation of the data rate in graphical form on a display. A fourth code is used to direct the outputting of a text representation of the one of the plurality of traffic classes on the display. The present invention has a variety of other codes to perform the methods described herein, and outside the present specification.

Col. 2:66-67 and 3:1-15. It is both unsurprising and irrelevant that Rakoshitz discloses that a software tool can perform separate functions using different code. Such is a probably a universal characteristic of software, in general. Again, the Second Answer misses the point. Claims 21 and 29 do not encompass all software that performs different functions. They are more specific than that.

Rakoshitz, for example, does not disclose “characterizing files stored in operable association with the individual servers according to type, using software operating on the individual servers,” as claims 21 and 29 define. The cited portion of Rakoshitz discloses “categorizing the data rate from the flow of information based upon at least one of a plurality of traffic classes,” but this is plainly not the same as characterizing files stored on individual servers. As discussed above, according to Rakoshitz, categorization of data takes place at a network gateway point for traffic flowing through the gateway, not for stored files. Characterizing stored files is neither disclosed by, nor inherent in the categorization process disclosed by Rakoshitz.

Issue 3: Rakoshitz does not disclose the claimed element of distributing a rule set to individual servers of a server group that defines rules for limiting serving of data from the individual servers depending on file type and a current state of the bandwidth usage.

On page 8 of the Second Answer, Rakoshitz at 8:44-55 is cited, which discloses in relevant part that “the present bandwidth management tool can be applied at . . . **an access to** a server farm (e.g., a group of servers located in a special part of a network close to an access link, e.g., in a web hosting environment)” (emphasis added.) Rakoshitz, consistent with its entire disclosure, here again discloses implementing a bandwidth tool at an access or gateway point. Under any reasonable claim construction, this differs from claims 21 and 29, which require “distributing a rule set to **individual servers** of the server group, wherein the rule set defines rules for limiting serving of data from the individual servers depending on file type and a current state of the bandwidth usage.” Under the argument set forth in the Answer, deploying a software tool “at an access to a server farm” is deemed to be the same as distributing the recited rule set “to individual servers of the server group.” This is neither reasonable, nor correct. Claims 21 and 29 clearly require the server group to comprise a file storage location that is distinct from the communication link where bandwidth use

is monitored. In contrast, Rakoshitz discloses implementing the entire tool at a single gateway or access point.

In addition, Rakoshitz does not disclose distributing a “rule set” that “defines rules for limiting serving of data from the individual servers depending on file type and a current state of the bandwidth usage,” as claims 21 and 29 require. Instead, Rakoshitz discloses:

The present tool identifies data flows at a network' site based on traffic classes. A traffic class is any combination of the following, but is not limited to these: IP address, subnet, network, netgroup, or range of source or destination; URL of the sender or group of URLs; Service (e.g., HTTP, FTP) or groups of services; FTP and HTTP, file types can be selected as well; Time of day, day of week/month; and Inbound and outbound information.

Col. 13:13-21. The key point here is that Rakoshitz discloses classifying data flows (as flowing through a gateway or access point), not files stored on file servers, as claimed. More exactly, Rakoshitz does not disclose the claimed element of distributing a rule set to individual servers of a server group, and instead teaches that traffic is controlled by managing data flows through a gateway point.

II. REJECTION UNDER 35 U.S.C. § 102(e) OF CLAIMS 24 & 32

Issue 1: Rakoshitz does not disclose crawling through a memory of an individual server to identify groups of files configured to be aggregated into a larger file.

The Second Answer does not effectively rebut the arguments set forth in the Appeal Brief demonstrating the deficiency of Rakoshitz with respect to claims 24 and 32. Instead, the Second Answer cites a portion of Rakoshitz (col. 2:66-67, 3:1-15, quoted above) in support of an argument that “Rakoshitz does disclose having the bandwidth controlling/profiling tool present in the memory of a network computing device.” Second Answer, p.9:17-19. Even if this is taken as true, the Answer does not point out where Rakoshitz shows identifying “associated groups of files, wherein each of the groups of files is configured to be aggregated into a larger file,” as claims 24 and 32 require. In fact, Rakoshitz does not disclose this element at all.

Moreover, the answer does not rebut the arguments on pages 12 and 13 of the Appeal Brief to the effect that Rakoshitz does not disclose “crawling” for files, as crawling is defined in the present application at page 10:28-29. Specifically, Rakoshitz does not disclose searching a file storage memory of a web server to classify files found there. Under any reasonable interpretation of “crawling” in view of the specification, merely operating a bandwidth management tool in a computer to classify files passing through a network gateway does not read on crawling through a file storage memory to identify files. Rakoshitz is not directed to bandwidth management at other than an access or gateway, and fails to disclose searching a file storage memory to identify stored files.

III. REJECTION UNDER 35 U.S.C. § 102(e) OF CLAIMS 25 & 33

Issue 1: Rakoshitz does not disclose crawling through a memory of storage device operable associated with the individual server to identify files that do not contain hyperlinks and are not identified by hyperlinks in other files stored by the storage device.

As discussed above in connection with claims 24 and 32, Rakoshitz does not disclose crawling through a memory to identify files. Claims 25 and 33 more explicitly define crawling through a “storage device,” so the deficiencies of Rakoshitz should be even more apparent. Rakoshitz merely discloses operating a bandwidth management tool in memory but nowhere discloses searching a storage device to identify files.

Moreover, Rakoshitz fails to disclose identifying “files that do not contain hyperlinks and are not identified by hyperlinks in other files stored by the storage device.” In rebuttal, the Second Answer cites col. 7:61-67 through col. 8:1-4, which discloses:

The present invention can also be used with a number of various files. For example, a number of common applications, such as FTP and HTTP, can handle a wide variety of files. The file types being transferred and downloaded place different demands on the underlying infrastructure. Index and HTML files take up limited bandwidth but have very mundane contents. On the other hand, GIF, JPEG and MPEG, RA and A VI files take up a lot more bandwidth but provide a rich multimedia experience to the end-user. In fact, push technologies such as PointCast basically download rich-multimedia bandwidth-intensive files.

This amounts to nothing more than a general disclosure that some files take up more bandwidth than others, as one might expect to see in a disclosure directed towards bandwidth management. However, neither here nor anywhere else does Rakoshitz disclose identifying files of the claimed type, “that do not contain hyperlinks,” etc. Claims 25 and 33 define crawling through file storage to identify a specific type of file type that is not otherwise described in Rakoshitz: i.e., “files that do not contain hyperlinks and are not identified by hyperlinks in other files stored by the storage device.” Such files may or may not comprise HTML or non-HTML files; for example, a non-HTML file may be linked to an HTML file and therefore would not be identified in a search for this type of file. On the other hand, the same type of non-HTML file would be identified if not named by a link in another stored file. Rakoshitz includes no indication that such files are of interest, nor of looking for files on the claimed storage device. Indeed, Rakoshitz fails to disclose any analysis of stored files, with or without reference to other stored files. Rakoshitz therefore does not disclose this element.

IV. REJECTION UNDER 35 U.S.C. § 102(e) OF CLAIMS 27 & 35

Issue 1: Rakoshitz does not disclose distributing a replacement rule set to individual servers of the server group in response to changes in bandwidth usage.

As discussed above in connection with claims 21 and 29 (Group 1), Rakoshitz fails to disclose or suggest distributing rule sets to individual servers of a server group,

as claimed. Instead, Rakoshitz discloses classifying data flows (as flowing through a gateway or access point), not files stored on file servers. Failing to disclose distribution of a rule set in the claimed manner, it follows naturally that Rakoshitz also fails to disclose distributing a replacement rule set in the claimed manner. As explained above and in the Appeal Brief, Rakoshitz does not disclose distributing rules to individual servers of a group.

However, it is argued on page 10 of the Second Answer that Rakoshitz discloses the claimed element at col. 9:24-62. Rakoshitz here discloses:

The bandwidth management tool 208 can be predominantly software based and is substantially free from any significant hardware or software changes in the network. In a preferred embodiment, the bandwidth management tool 208 can be loaded onto a server without any changes to hardware. In an alternative preferred embodiment, the tool can install, configure, and operate on a conventional IBM compatible PC running an operating system such as, for example, Windows NT, but can be others. The tool can be deployed at any appropriate point in the network data path. The tool can also be stand-alone at the WAN access point (e.g., behind the Internet access router or behind a firewall), with a conventional firewall or with an NT based proxy/caching server or application server (e.g., a Web server).

Tool 208 performs incoming and/or outgoing management of information over the network of computers. In a specific embodiment, traffic management tool 208 performs inbound and outbound monitoring and control of flows by application, source address, destination address, URL, time of day, day of week, day of month, and other variations. In a specific embodiment, tool 208 also monitors, controls, and produces reports and alarms, which can enhance a whole spectrum of traffic monitoring and control activities ranging from bandwidth/latency control to capacity planning.

In a specific embodiment, the bandwidth management tool adapts to "real" changes on any pre-existing networking system. For example,

network infrastructure management involves a continuous process of monitoring, reporting, and deploying changes to match network growth or changing needs in a growing office, for example. These changes exist at various levels and time scales. As merely examples, the network changes, can be to enforce a QoS Policy for a critical service, add WAN bandwidth, segment the network, upgrade a router, choose a guaranteed service level for a web site (e.g., user's own wet site), or notify "Mr. Hog" (i.e., a user occupying too much bandwidth) that he should schedule his large personal downloads at more prudent times such as late at night, for example.

Rakoshitz here discloses that a bandwidth management tool can “adapt[s] to ‘real’ changes on any pre-existing network system.” Clearly, Rakoshitz lacks any specific disclosure of changing a rule set for “limiting serving of data from the individual servers depending on file type and a current state of bandwidth usage,” as claims 27 and 35 define. Instead, the examples given by Rakoshitz and cited above concern systemic changes to the network, e.g., adding WAN bandwidth, etc., not transient conditions such as the claimed “current state of bandwidth usage.”

Furthermore, irrespective of the type of network changes for which adaptations are made, Rakoshitz nowhere discloses controlling bandwidth usage by distributing rule sets to individual servers. Instead, as explained above in Section I regarding the Group I claims, Rakoshitz discloses controlling flow at access points and gateways to prioritize data flow downstream of servers that are serving the files. Controlling flow at access points and gateways cannot reasonably be deemed to read on distributing rules sets to individual servers. In addition, the Group IV claims add the additional limitation that

replacement rule sets are distributed in response to changes in bandwidth usage, which is also not disclosed by Rakoshitz.

Issue 2: Rakoshitz does not disclose distributing cooperative elements at both communication links and file storage locations, as claimed.

The Second Answer cites Rakoshitz at col. 8:44-55, which, as discussed above, merely provides examples of different network locations at which a bandwidth management tool may be deployed. It has not been demonstrated that Rakoshitz discloses the specific distribution of cooperative elements defined by claims 27 and 35. Merely disclosing that the bandwidth tool may be deployed at multiple locations does not, under any reasonable claim construction, disclose the specific step or system of

“distributing a replacement rule set to individual servers of the server group when the current state of the bandwidth usage changes by more than a specified amount, wherein the replacement rule set replaces the rule set and defines rules for limiting serving of data from the individual servers depending on the file type and a current state of bandwidth usage,”

as claims 27 and 35 recite. These claims include numerous limitations that are not encompassed by a general disclosure that a bandwidth management tool may be deployed at different network points. The bare disclosure of Rakoshitz lacks any disclosure of a replacement rule set. It lacks any disclosure of providing the replacement rule set to individual servers of a server group. And it lacks any disclosure

of a rule set that defines rules for limiting serving of data from the individual servers depending on the file type and a current state of bandwidth usage.

As explained above under Section 1, Issue 1, Rakoshitz merely discloses that a bandwidth management tool may be implemented at various access points. Indeed, this is the only practical place to install a tool such as Rakoshitz that operates on the principle of analyzing and controlling network traffic. The Board is referred to Section 1 above for a more complete discussion of the deficiencies of Rakoshitz in this regard.

V. CONCLUSION

The arguments set forth in the Appeal Brief have not been successfully rebutted by the Second Answer, as shown by the discussion above. Appellants respectfully request the reversal of the rejection of currently pending Claims 21-36, and allowance of these claims forthwith, for the reasons set forth in the Appeal Brief and above.

Respectfully submitted,

Date: April 9, 2007

/Jonathan Jaech/
Jonathan Jaech
Attorney for Applicant
Registration No. 41,091

Serial No. 09/932,431
April 9, 2007
Page 16

CUSTOMER
NUMBER

58688

PATENT TRADEMARK OFFICE

Connolly Bove Lodge & Hutz LLP
355 South Grand Avenue
Suite 3150
Los Angeles, CA 90071-1560
(213) 787-2500